

Autonomous Multi-sensor Coordination: The Science Goal Monitor

Anuradha Koratkar^a, Sandy Grosvenor^b, John Jung^c, Jenny Geiger^c

Goddard Earth Science and Technology Center, University of Maryland Baltimore County,
Baltimore MD USA 21250

Science Systems and Applications, Inc., NASA/GSFC, Greenbelt, MD USA 20771
Advanced Architectures and Automation Branch, NASA/GSFC, Greenbelt, MD USA 20771

ABSTRACT

Next-generation science and exploration systems will employ new observation strategies that will use multiple sensors in a dynamic environment to provide high quality monitoring, self-consistent analyses and informed decision making. The Science Goal Monitor (SGM) is a prototype software tool being developed to explore the nature of automation necessary to enable dynamic observing of earth phenomenon. The tools being developed in SGM improve our ability to autonomously monitor multiple independent sensors and coordinate reactions to better observe the dynamic phenomena. The SGM system enables users to specify events of interest and how to react when an event is detected. The system monitors streams of data to identify occurrences of key events previously specified by the scientist/user. When an event occurs, the system autonomously coordinates the execution of the users' desired reactions between different sensors. The information can be used to rapidly respond to a variety of fast temporal events. Investigators will no longer have to rely on after-the-fact data analysis to determine what happened.

This paper describes a series of prototype demonstrations that we have developed using SGM and NASA's Earth Observing-1 (EO-1) satellite and Earth Observing Systems' Aqua/Terra spacecrafts' MODIS instrument. Our demonstrations show the promise of coordinating data from different sources, analyzing the data for a relevant event, autonomously updating and rapidly obtaining a follow-on relevant image. SGM is being used to investigate forest fires, floods and volcanic eruptions. We are now identifying new earth science scenarios that will have more complex SGM reasoning. By developing and testing a prototype in an operational environment, we are also establishing and gathering metrics to gauge the success of automating science campaigns.

Keywords: remote sensing, onboard analysis, spacecraft autonomy, sensor webs, optimized observations

1. INTRODUCTION

Many dramatic earth phenomena are dynamic and coupled. In order to fully understand them, we need to monitor them and obtain timely coordinated multi-sensor observations from widely dispersed instruments. The need for dynamic coordinated multi-sensor observations has given rise to the concept of a sensor web, which characterizes a future Earth observing system concept far more capable than today's observing systems. Sensor webs or autonomous collaborative systems will monitor the intrinsically dynamic behavior of a wide variety of naturally occurring (e.g., wild fires, flash floods, hurricanes, and volcanoes) and human-induced (e.g., toxic spills, pollution) events and phenomena. It is envisioned that sensor webs will:

- Maximize the return of only the most useful scientific information;
- Minimize overall response time of the system when monitoring rapidly evolving phenomena;
- Conduct near-real-time synthesis or "fusion" of information from multiple assets. Numerical models will also be considered as assets and will contribute toward improving our understanding of the complex, interrelated processes that drive the formation and evolution of environmental phenomena.

Unlike today's "stove-piped" missions where scientific goals are often achieved using passive observing strategies, the dynamic nature of the future observing systems implies that these goals can be achieved using reactive, and proactive dynamic observational strategies that use all available interconnected resources (e.g., various space craft, balloons, numerical models and other distributed/configurable sensors).

We are developing a prototype software architecture called the Science Goal Monitor (SGM; <http://aaa.gsfc.nasa.gov/SGM>) to investigate the potential benefits of the sensor web, determine the properties that it should possess, and determine the best methods with which to develop, integrate, deploy and then operationally use the sensor web's technologies and information systems,. SGM is being developed for time-variable phenomenon where goals can be defined as measurable objectives with contingency plans for follow-on work. Hence, SGM is designed to be adaptable to many different science domains that require rapid response. For example, SGM will recognize fast temporal space science events such as gamma ray bursts or earth based evolving events such as forest fires, floods and volcanic eruptions. Therefore, our project has both space and earth science collaborations. In the space science domain, we are implementing a prototype for dynamic automated reactions to intrinsically varying astronomical phenomenon using one of the Small and Moderate Aperture Research Telescope System (SMARTS) telescopes. Our primary objectives with the SGM/SMARTS collaboration is to prototype a small set of test scenarios in an environment that is scientifically and operationally realistic, and one in which we can measure the effectiveness of SGM. In the earth science domain we are collaborating with the Earth Observing-1 (EO-1) team to evaluate how multiple sensors consisting of EO-1/Terra/Aqua satellites can react dynamically to obtain rapid observations of evolving earth science events. The SGM prototype will help us evaluate the sensor web observing strategy and understand how the science community will operationally use the sensor web.

Introduction of flexible scheduling and autonomously reacting to science driven events implies a certain amount of automation. There are a number of challenges inherent in infusing autonomy, especially in an existing environment that was not built for autonomy. The SGM project is a proof-of-concept effort to address these challenges. A science goal monitoring system is an enabler for spacecraft autonomy and an important component of the sensor web concept.

2. SCIENCE GOAL MONITOR

2.1 What is the Science Goal Monitor (SGM)

In order to capture transient/dynamic events, a sensor web system must have an automated reactive capability built into the sensor web operations. To implement such a dynamic data collection strategy, *we must teach our sensors to intelligently achieve the user/scientists' goals*. At a minimum we see that sensor webs must include the ability to:

- Schedule flexibly and react autonomously to science/user driven events.
- Understand higher-level goals of a scientist/user defined campaign.
- Coordinate various space-based and ground-based resources/sensors effectively and efficiently to achieve goals.

The SGM system is being designed to understand high-level user goals, translate them into a machine interpretable format, monitor the relevant incoming data stream for events related to these goals, autonomously recognize and react to the events in a timely fashion, and autonomously schedule the necessary observations to meet the user goals. When an event occurs, the system will instruct the scheduler to reprioritize observations/simulations based on rules defined in the science/user program/campaign. The data stream can be either from the relevant hardware or from numerical models. SGM provides users with visual tools to capture user goals in terms of measurable objectives. The users specify what to look for and how to react in descriptive rather than technical terms. Our prototype is designed for use in a distributed environment where some analysis might be performed onboard a spacecraft, while other analyses might be performed on the ground.

2.2 EO-1/Terra/Aqua Sensor Web

Using NASA's Earth Observing-1 (EO-1) satellite (<http://eo1.gsfc.nasa.gov>) and Earth Observing Systems' Aqua/Terra spacecrafts' MODIS instrument (<http://modis.gsfc.nasa.gov>), we have recently completed a series of prototype demonstration tests of varied evolving phenomenon such as: forest fires, volcanoes, floods and lake freezing. The EO-1/Terra/Aqua sensor web scenario using the SGM project demonstrates the use of multiple satellites in generating a simple sensor web to observe earth science phenomenon. In these demonstrations SGM analyzed data received from the MODIS instruments with a dynamic autonomous request for higher-resolution images from the EO-1 satellite based on a set of scientific criteria. Details of the various demonstrations so far are discussed in Section 3.

To discuss the various steps in the observing strategy we will consider a simple example related to forest fires. Figure 1 illustrates the data flow for the EO-1 forest fire demonstration.

- The scientist defines the high-level goal using the SGM. For example a high level goal could be, image the most recent significant fire at a higher resolution. This goal would have to be further dissected because the terms “recent”, “significant” and “higher resolution” are user dependent and need to be defined using additional parameters. The underlying goals of the observation/project are captured in a machine readable/understandable format. Once these details are known, *SGM autonomously recognizes and reacts in a timely fashion when goals are met*. For example, SGM can determine where to get the data from and which spacecraft are available for higher resolution imaging within the specified time period etc. The user need not generate such details. But, in the present prototype the user needs to identify the EO-1 instrument that must be used for the follow-on observation.
- SGM requests forest fire data from the MODIS Instrument center which e-mails alerts within 3 hours of image acquisition with fire lat/long information. The fire alert data are generated from images taken by Aqua and Terra with a 1km resolution per pixel.
- SGM determines the “most recent significant fire”, by correlating the MODIS Rapid Fire data with the information that is available at the Remote Sensing Applications Center in Utah (<http://www.fs.fed.us/eng/rsac>). These two data sets are used to refine the lat/long information for the fire of interest.
- SGM then automatically coordinates with the EO-1 Mission Operation Center to obtain a 30m higher resolution

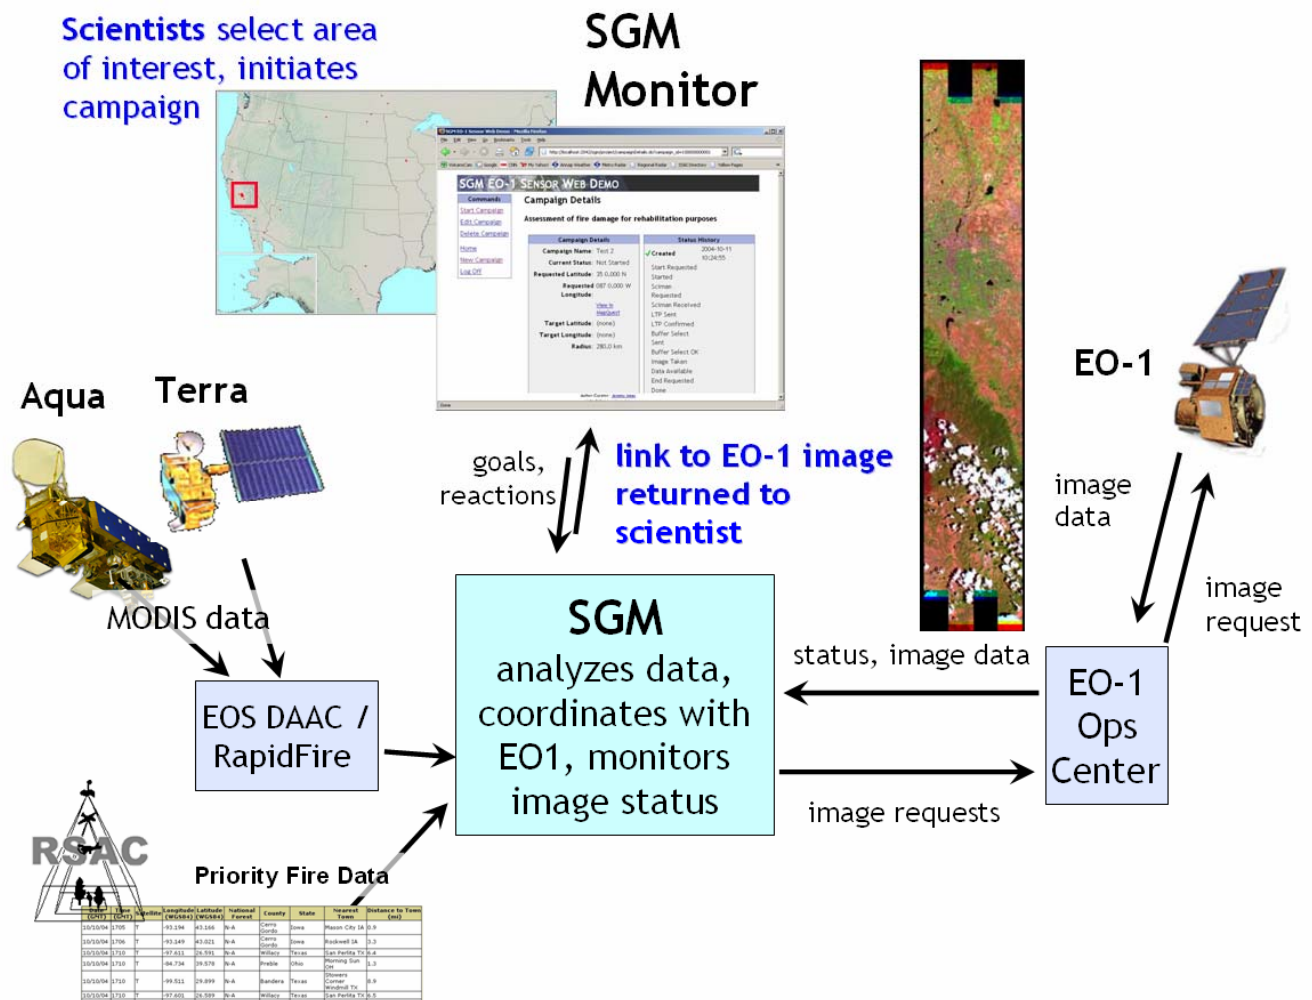


Figure 1: SGM data flow for the EO-1 forest fire demonstration

image with the user specified EO-1 instrument.

- After EO-1 notifies SGM that an observation is possible for the given location, EO-1 autonomously re-plans to take a closer look at the desired region. The EO-1 data are down linked to the UGS data center.
- SGM monitors the status of the program. Once data are obtained at UGS, it notifies the scientist and provides links to the science images (MODIS and EO-1) located at different locations.

3. SGM PROTOTYPE DEMONSTRATIONS

Our demonstrations thus far have focused on the basic capabilities of SGM. They show the promise of coordinating data from different sources, analyzing the data for a scientifically relevant event, and autonomously updating and rapidly obtaining a follow-on scientifically relevant image in a number of different science domains. The following prototype demonstrations have been conducted and are being used to date.

3.1 Forest Fires

A number of space-based assets are currently being used for detecting active fires, mapping burned area, assessing fire susceptibility and estimating fire emissions. Burned areas are re-observed over time to monitor the recovery of the area. The sensor web concept is ideally suited for fire management because timely delivery of meteorological and satellite data at the appropriate spatial scale is essential for many of the activities.

High-resolution data are often needed to provide information at a finer spatial scale such as assessing fire damage and monitoring post fire recovery. Therefore, for our first prototype demonstration we choose to obtain a high-resolution image of a fire for post-fire damage assessment. In this scenario, SGM determined the largest significant fire in a certain region, and then directed EO-1 to take a higher resolution image. In this scenario, SGM served both as a science

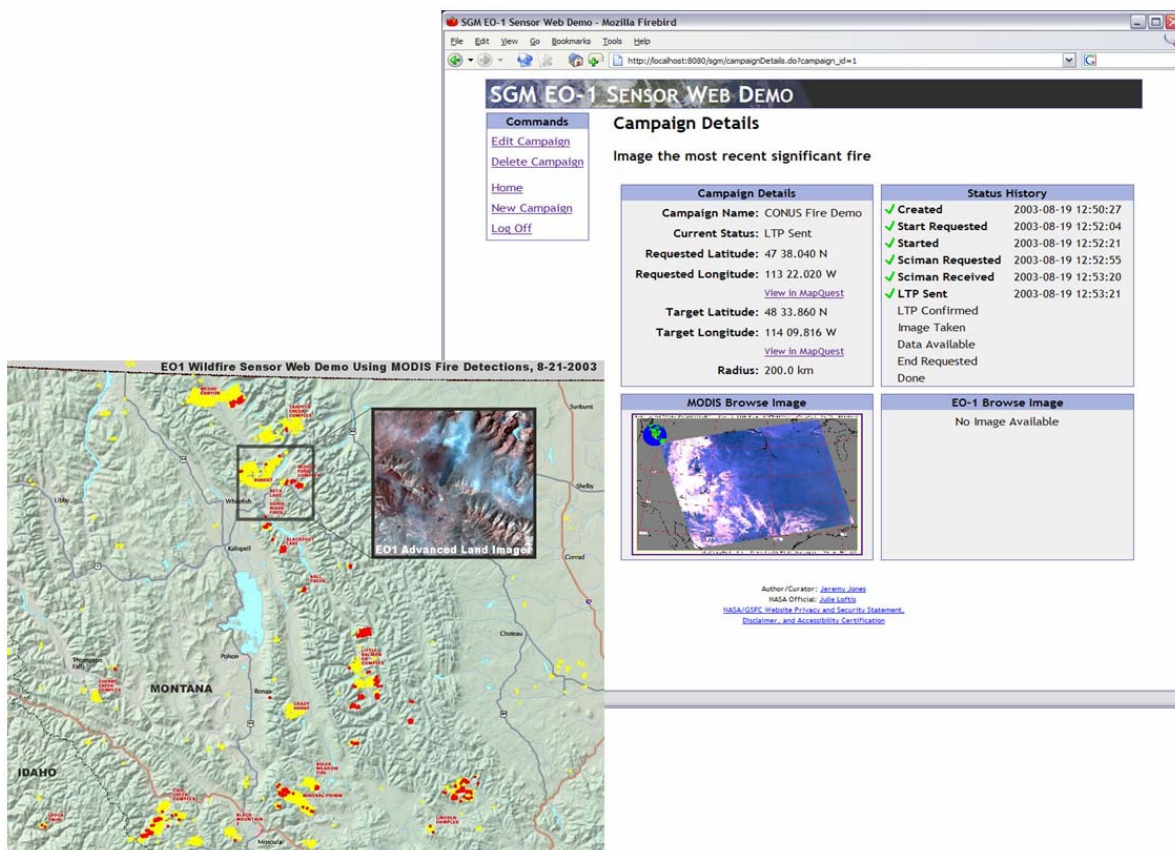


Figure 2: Details of the SGM web interface

analyzer and a multi-mission coordinator. The significant steps for this scenario are:

- The scientist currently enters their region of interest by entering a latitude, longitude, and radius.
- SGM monitors the daily list of active priority fires from the Remote Sensing Applications Center in Utah (<http://www.fs.fed.us/eng/rsac>), and identifies priority fires from the RSAC that are located within the region of interest.
- SGM analyzes the recent history of the fires from the MODIS Rapid Fire data in the region of interest to isolate the latest center of activity.
- Using SGM's updated "centroid" of the fire, SGM coordinates with the EO-1 planning systems to request a high-priority high-resolution image of the fire and monitor the status of the request.
- SGM's web-based user interface provides the user with a live display of the status of his/her image request and automatically links to the new EO-1 image when it became available.

The SGM coordination and analysis provided new data to the US Forestry Service within 48 hours, compared to a typical lead-time of up to 14 days for preplanned observations. The SGM and EO-1 teams have performed several demonstrations over the past year, utilizing fires in the southern hemisphere and early spring events in Florida, Southern California and the Southwest U.S.

3.2 Volcanoes

Over 100 volcanoes erupt somewhere on earth every year, some with devastating consequences. Currently, only a limited number of volcanoes are being monitored. When there is a significant event, manual intervention to acquire relevant data is the standard observing strategy. An observing strategy where scientists can "catch" the eruption is extremely useful. A sensor web can provide the means to effectively and efficiently monitor a large number of volcanoes and then obtain timely data to observe an eruption.

In the volcano scenario we focused on demonstrating SGM's ability to conduct long term monitoring of multiple targets, and then conduct follow-on observations when an eruption is detected. The volcano demonstrations show the feasibility of quickly gathering high-resolution satellite images when an event is detected from one of several independent, geographically remote sensors and with little human involvement. The significant steps for our initial volcano scenario are:

- Scientist specifies a prioritized list of volcanoes to monitor for new eruptions.
- SGM monitors each volcano site again using Rapid Fire datasets from the MODIS instruments.
- When "hot pixels" above the scientist's specified temperature threshold are detected on a volcano, SGM deduces that an eruption is in progress. Then SGM automatically coordinates with EO-1 to request a high-resolution image of the volcano area. If more than one eruption is detected, SGM selects the highest priority site for the follow-on EO-1 observation.
- The user receives the high-resolution image via the SGM web front-end.

For the first volcano demonstration SGM relied solely on Rapid Fire data to deduce an eruption. We have now integrated in-situ monitoring information from tilt meters into SGM. A follow-on EO-1 observation is requested by SGM when either the tilt meter data indicate a seismic event or a temperature threshold from Rapid Fire data is encountered.

3.3 Floods

Floods affect large areas of the Earth. They are not reliably predictable. Hydrological data from in situ sensors is very sparse. Also in-situ sensors cannot provide information on the geographic extent of the flood. Sensor web observational strategies can be used to map, measure, and monitor floods. Similar to the fire scenarios, SGM can be used to automate an on-going monitoring program for a flood area, obtaining images during flood, and then monitoring post-flood recovery progress.

To demonstrate SGM's applicability to monitor river flooding, we interfaced SGM to data from the Dartmouth Flood Observatory (<http://www.dartmouth.edu/~floods>), which monitors various rivers and wetlands around the world using the QuickSCAT scatterometer instrument. SGM monitored this data for flooding alerts concerning a user-specified river. For our demonstration, the Brahmaputra River in India was selected. SGM detected an alert and, as in the fire and volcanoes scenarios, sent a request for a high-resolution image acquisition to EO-1 based on the latitude and

longitude specified in the flood alert record. Future flood scenarios will use ground in-situ sensors to predict potential flooding before it occurs, which will drive subsequent EO-1 observations of the target area similar to the volcano scenario.

3.4 Lake Freezing

The University of Wisconsin maintains a series of buoys in Sparkling Lake that measure surface water temperature. The goal of this scenario is to monitor the data from those buoys to determine when the lake's first freezing occurred, then to take an image of the lake area as soon as possible to characterize the lake environment during the time of transition. The science goal is to determine the boundaries between liquid and solid water. The EO-1 team is using these observations to refine algorithms being developed and tested for on-board detection of water and ice. These algorithms will be valuable for future autonomous planetary exploration. SGM monitored the buoy readings for several days and triggered an EO-1 observation as soon as the temperature readings showed the lake's surface was beginning to freeze.

3.5 Integration of Weather Forecast to Prioritize the Satellite Tasking

Often many of the detectors on earth observing satellites observe cloud covered tracks and ship the data down. For many detectors this is useless. To obtain useful scientific data these observations must be cloud free. With the increase of onboard processing power, an observation can be analyzed on the satellite itself to determine if the observation is cloud free before it is transmitted back to earth. In the sensor web domain, we can determine "near real time" if an observation will have a high probability of being cloud free by accessing real-time weather data. If the observation has a high probability of being cloud free, then the observation is executed; else the satellite time can be used more effectively for another observation. In this manner the sensor web can be used to maximize the returns from our assets.

Step 1:
Submit 2 competing
acquisition targets

Step 2: Automatically
query GOES RT cloud data
1-4 hours before overflight

Step 3: Autonomously
command EO-1 to
switch to alternate 2
target


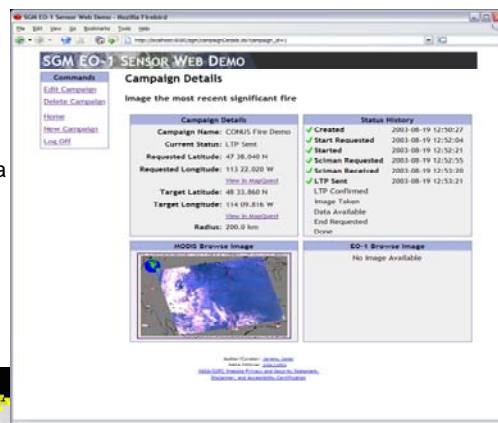
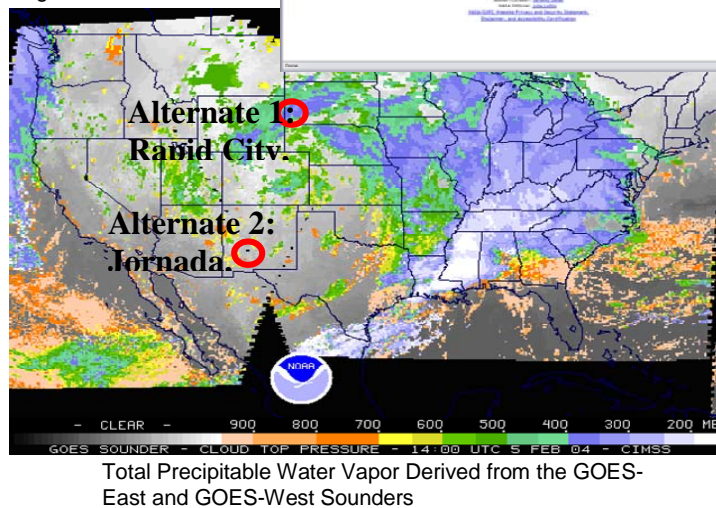


Figure 3, SGM data flow for the EO-1/GOES cloud cover demonstration

To demonstrate that cloud free images to maximize the scientific value of the image can be obtained, SGM accessed real-time weather data from the GOES satellite, and then coordinated with EO-1 planning. Using a mix of the scenarios described above, SGM scheduled requests for primary and alternate scenes to be taken aboard EO-1. One to four hours before EO-1 image was scheduled to be executed, SGM accessed and analyzed the latest GOES cloud cover data and autonomously notified EO-1 if the primary scene was cloud covered, in which case, EO-1 took an image of the alternate scene. This allows EO-1 to best use its high demand imaging time while minimizing lost imaging time to cloud-covered scenes (see Figure 3).

3.5 Multi-phenomenon Monitoring

Our latest demonstrations have begun to mix the above scenarios using scientist-specified priorities so that SGM can monitor several emerging events, compare the priorities of the requested images, ensure that the scene is cloud-free, and then direct EO-1 to pick the highest priority scenes. This scenario will allow us to understand how to handle competing priorities in an environment where there are multiple observing alternatives.

4. SGM STATUS AND FUTURE PLANS

While each of the scenarios described earlier has been successfully demonstrated, we continue to operate in a demonstration, prototype mode. As such, SGM runs only for short periods for pre-planned demonstration windows. It is not yet analyzing data in a continuous fashion and requesting scenes whenever events occur. We hope to extend the SGM's capabilities to be able to continually monitor disparate sensors and integrate image requests into EO-1's production schedule on an on-going basis. We also hope to expand the SGM and EO-1 link to perform science campaigns with more complex reasoning and across a wider variety of remote sensors and demonstrate truly autonomous science driven monitoring. We do not yet have an automated ability for SGM to access EO-1's imaging schedule, commitments and priorities. With access to such data, we could expand the multi-target demonstration to monitor potential targets over time, and match the priority of desired images with the priority of the scheduled images. SGM driven high-priority image requests might override lower priority pre-scheduled images but not be able to override other higher priority images. The full potential of the SGM and EO-1 collaboration could then be seen, when SGM autonomously discovers and coordinates an image of a scientific event so that the scientists/users have the images on-hand during the critical decision making periods. Finally, we plan to improve the end-user's (scientists) interface so that new scenarios and campaigns are easier to create and manage.

5. LESSONS LEARNED

Legacy spacecraft can be a part of a sensor web. In our demonstrations we have shown that "sensor web" like automated observing scenarios can be achieved with assets available today to the community. But there are a number of automation challenges that exist because the environments were not built with autonomy in mind. We discuss some of the challenges that we had to overcome.

5.1 Automation of manual tasks in an existing environment

NASA mission software has historically followed very conservative development and implementation methodologies; hence there are a number of routine tasks that require manual intervention. Further, each mission typically includes various bits of custom software to do different things, which is very labor-intensive to automate given that there are few standard interfaces to these tools. To change an existing stable (albeit heavily manual) system into a completely autonomous system is therefore not realistic with the limited funds available. However, by looking at the full spectrum of automation possibilities, significant progress can be achieved by identifying areas where a manual task can be automated by relatively simple means such as providing automated front-end or "wrapper" around the manual front-end for tasks that are repetitive in nature. Working closely with the existing mission engineering staff is essential since they are the ones who know the mission operations software and are responsible for the ongoing stability of software.

With the EO-1 mission, we have worked with the EO-1 engineers to use the legacy nature of EO-1 to both our advantages. EO-1 was not designed or built with any expectation of automated operations; consequently, we have run into barriers, that limited the kinds of interactions we could achieve in the SGM software. But the EO-1 team's

approach has been to work with the SGM team to add some automation, thereby enabling links to SGM, while applying automation that helps reduce the level of manual effort for the EO-1 operations team.

5.2 Changes in mission operation

One desired outcome of detecting temporal events is the modification of the observing plan/schedule for the spacecraft. However, legacy missions' operations are not flexible enough to accommodate unexpected changes while maintaining a consistent plan and schedule. Satellites such as EO-1 simply do not have sufficient onboard processing to even entertain the idea of onboard planning. As onboard processing improves rapidly, the barrier to on-board automation is not technology-driven, but the conservative mindset, (developed over the years), that the plan must be manually reviewed prior to execution. In order to overcome this barrier, successful test systems must be developed and proven before a change in operations strategy will occur.

EO-1's schedule is determined through a highly manual process and is finalized 24-48 hours in advance and shipped up to the satellite as a nearly static schedule. This significantly limits how reactive a system can be, but it has not prevented implementing some dynamic adjustments. For example, in our ground-based cloud scenario, the EO-1 schedule is actually built with commands to execute both the primary and alternative scenes. EO-1 has a limited ability to by-pass scenes on a short time-frame, which SGM uses to request which scene to actually execute.

The response time of any automated system built on an existing legacy system is largely bound by the response time of the legacy system. If the event is short-lived and the operations system cannot dynamically implement a change in the schedule fast enough, then we could miss capturing the event entirely. This can be a major problem if rapid changes are desired. Even though the legacy software may be able to handle more flexible scheduling, if the mission operation policies do not allow for that level of flexibility, it will not be possible. Responding to very rapid events will involve a migration of science and scheduling responsibility onto the spacecraft. Both SGM and EO-1 are working on means to explore the best balance between ground-based responses and on-orbit responses.

CONCLUSIONS

Spacecraft will be an integral part of the future sensor webs, and they will need to have some flexibility in operations if the sensor web has to dynamically respond to events. Developing a spacecraft with flexible scheduling, which can autonomously react to unpredictable events is not just a leap forward in automation, but a large change in operations paradigm. There are a number of challenges that need to be overcome to change the present NASA mission operation strategies.

Our SGM demonstrations so far show that

- Relevant science events can be accurately detected and that the details of the follow-on observations can be correctly captured and transmitted to the appropriate location autonomously.
- Existing assets that were not built for autonomy or flexible scheduling, and that have limited onboard processing, can still be autonomously coordinated and dynamically scheduled. Automation of a legacy system is easier to achieve if that automation is layered on top of the existing system. However, the limit of such automation is directly related to the flexibility of the existing system design.

In SGM we are developing an interactive distributed system that will use onboard processing and storage combined with event-driven interfaces with ground-based processing and operations, to enable fast reaction to time-variable phenomena. We are currently developing prototypes and evaluating the effectiveness of the system. Although we have not completed our project, we have had some success. Our dynamic science analysis and autonomous multi-sensor coordination has been highly appreciated scientists are looking forward to attempting more complex problems.

Acknowledgements: This work is funded by NASA's Code R under the Computing, Information and Communications Technologies (CICT) program.